Optimization of distributions differences for classification

Mohammad Reza Bonyadi, Quang M. Tieng, David C. Reutens,

Abstract—In this paper we introduce a new classification algorithm called Optimization of Distributions Differences (ODD). The algorithm aims to find a transformation from the feature space to a new space where the instances in the same class are as close as possible to one another while the gravity centers of these classes are as far as possible from one another. This aim is formulated as a multiobjective optimization problem that is solved by a hybrid of an evolutionary strategy and the Quasi-Newton method. The choice of the transformation function is flexible and could be any continuous space function. We experiment with a linear and a non-linear transformation in this paper. We show that the algorithm can outperform 6 other state-of-the-art classification methods, namely naive Bayes, support vector machines, linear discriminant analysis, multi-layer perceptrons, decision trees, and k-nearest neighbors, in 12 standard classification datasets. Our results show that the method is less sensitive to the imbalanced number of instances comparing to these methods. We also show that ODD maintains its performance better than other classification methods in these datasets, hence, offers a better generalization ability.

arXiv:1703.00989v1 [cs.LG] 2 Mar 2017

1 Introduction

The ultimate goal of a supervised classification method is to identify to which class a given instance belongs based on a given set of correctly labeled instances. A classifier in this paper is defined [1] as follows:

Definition 1. *(Discriminative classifier) Let $S, S_1, S_2, ..., S_c$ be sets of instances, $|S_i| = m_i$, $S_i \cap S_j = \emptyset$ for all $i, j \in \{1, ..., c\}$, $i \neq j$, $\cup_{i=1}^{c} S_i = S$. A classifier $\psi_\kappa(x)$, $\psi_\kappa : S \rightarrow \mathbb{R}$, aims to guarantee*

$$\forall i \in \{1, ..., c\}, \forall x \in S_i, P(\psi_\kappa(x) = i|x) = 1.$$

where κ is a set of configurations for the procedure $\psi_\kappa(x)$, and P is the probability measure.

In this paper, we consider a special case of classification problems where all members of S are in $\mathbb{R}^n$ (so called feature space). We also assume that feasible values for x^i (called a variable throughout this paper), the i^{th} element of the instance $\vec{x}$, are ordered by the operator "$\leq$" (i.e., x^i is not categorical).

In reality, only a subset of S is given (the training set) for which the classes are known. It is then critical to find the best κ in a way that $\psi_\kappa(\vec{x})$ is the true class of any $\vec{x}$ in the training set and, ideally, all possible instances in S. However, the distribution of instances in each class S_i is unknown, making the best estimation of κ challenging.

Classification is required in many real-world problems. Although there have been many classification methods proposed to date, such as multi-layer perceptrons [2], decision trees, support vector machines [3], and extreme learning machines [4], there are still some limitations associated with many of these methods. Some methods are sensitive to imbalanced number of instances in each class [5]. Also, non-linear classification methods outperform linear classification methods in the training set, however, may end up having worse performance when they are applied to new instances, an issue known as overfitting [6].

We propose a new classification algorithm called optimization of distributions differences (ODD) in this paper. ODD aims to optimize distribution of instances in different classes to ensure they do not overlap. ODD finds a transformation, $F : \mathbb{R}^n \rightarrow \mathbb{R}^p$, n is the number of dimensions (features) for each instance and p is a positive integer, in a way that the distance between the gravity centers of the instances in different classes is maximized while the spread of the instances within the same class is minimized. If such transformation exists, the instances could be assigned to each class based on their distances to the centers of the classes. We formulate the optimization of this transformation as a multiobjective optimization problem with two sets of objectives. The first set of objectives ensure that the centers of different classes are as far as possible from one another by defining $c(c-1)/2$ objectives, c is the number of classes. The second set of objectives ensure that the spread of instances within each class is minimized, defined by the norm of the eigenvalues of the covariance matrix of instances in each class, that adds c extra objectives to the system. We solve this problem using a combination of evolutionary algorithms [7] and conjugate gradient methods [8]. We experiment with linear and non-linear transformations and show that the method can outperform existing classification methods when they are applied to 12 standard classification benchmark problems and 4 artificial classification problems. We also show that the algorithm is not sensitive to the imbalance number of instances in the classes, assuming that the given instances for training represent the classes distribution parameters to some extent. Our experiments indicate that the method outperforms both non-linear and linear classifiers in terms of generalization ability.

We structure the paper as follows: section 2 outlines some background information on classification methods

All authors are with the Centre for Advanced Imaging (CAI), the University of Queensland, Brisbane, QLD 4072, Australia. M. R. Bonyadi (reza@cai.uq.edu.au, rezabny@gmail.com) is also with the Optimisation and Logistics Group, The University of Adelaide, Adelaide 5005, Australia.

and optimization algorithms we will use in this paper. Section 3 provides details of our proposed method, including the model, definitions, and optimization. Section 4 reports and discusses comparative results among the multiple classification methods on 12 standard benchmark classification problems. Sensitivity to imbalance number of instances in each class and overfitting are also discussed in that section. Section 5 concludes the paper and discusses potential future directions.

2 BACKGROUND

This section provides some background information about existing classification and optimization methods.

2.1 Classification methods

In this section, we describe classification methods that we use for comparison purposes.

2.1.1 K-nearest neighbor (Knn)

K-nearest neighbor (Knn) [9] works based on the assumption that instances of each class are surrounded mostly by the instances from the same class. Hence, given a set of training instances in the feature space and a scalar k, a given unlabeled instance is classified by assigning the label which is most frequent among the k training samples nearest to that instance. Among many different measures used for distance between instances, Euclidean distance is the most frequently used for this purpose.

2.1.2 Naive Bayes (NBY)

NBY is a classification algorithm that works based on the Bayes theorem [10]. The aim of the algorithm is to find the probability that a give instance $\vec{x}$ belongs to a class c, i.e., $P(c|\vec{x})$. To calculate this, NBY uses the Bayes theorem as $P(c|\vec{x}) = \frac{P(\vec{x}|c)P(c)}{P(\vec{x})}$. The value of $P(c)$, $P(\vec{x})$, and $P(\vec{x}|c)$ can be all estimated from the given instances in the training set. As the instance $\vec{x}$ is in fact a vector that contains multiple variables, $P(\vec{x}|c)$ is estimated by $P(x_1|c) \times P(x_2|c) \times ... \times P(x_n|c)$ that indeed ignores the dependency among variables, a "naive" assumption.

2.1.3 Support vector machines (SVM)

The aim of SVM [3] is to find a hyperplane defined by the normal vector $\vec{\omega}$ that could separate two class of instances [3]. The separation is determined by the sign of $\vec{\omega}\vec{x}^T + r$ that indicates to which side of the hyperplane the instance $\vec{x}$ belongs. In other words, given a set of instances and their classes (supervised learning), the algorithm outputs an optimal hyperplane which categorizes instances.

One way to extend this algorithm to deal with multiclass classification problems is to use one-vs-all or one-vs-one strategies proposed in [11].

2.1.4 Multi-layer percenptrons (MLP)

MLP aims to optimize the parameters of a mapping from a set of input instances to their provided outputs to estimate the Bayes optimal discriminant [12]. The mapping can be linear or non-linear and could be presented in multiple layers. The algorithm minimizes the mean square of error between the generated outputs and expected outputs for each instance. One of the most frequently used optimization methods in MLPs is the Levenberg-Marquardt [13], [14], that is also used in this paper. See [15] for more details.

2.1.5 Linear discriminant analysis (LDA)

The aim of LDA is to calculate $\vec{\omega}$ for which $\vec{\omega}\vec{x}^T > k$ if the instance $\vec{x}$ belongs to the second class. Assuming that the conditional probability $p(\vec{x}|y = 0)$ and $p(\vec{x}|y = 1)$ (y is the label of $\vec{x}$) are both normally distributed with mean and covariance parameters $(\vec{\mu}_1, \Sigma_1)$ and $(\vec{\mu}_2, \Sigma_2)$, Fisher [16] proved that $\vec{\omega} = (\Sigma_1 + \Sigma_2)^{-1}(\vec{\mu}_2 - \vec{\mu}_1)$ and $k = \frac{1}{2}\vec{\mu}_2^T\Sigma_2^{-1}\vec{\mu}_2 - \frac{1}{2}\vec{\mu}_1^T\Sigma_1^{-1}\vec{\mu}_1$ could distinguish between the two classes. Considering $S_W = (\Sigma_1 + \Sigma_2)$ as a measure for the within-class spread and $S_B = (\vec{\mu}_2 - \vec{\mu}_1)(\vec{\mu}_2 - \vec{\mu}_1)^T$ (T is the transpose operator) as a measure for between-class spread, Fisher's value for $\vec{\omega}$ ensures that $\frac{W^T S_B W}{W^T S_W W}$ is maximized. $\vec{\omega}$ is the norm of a hyperplane that discriminates the two classes and k is the shift to ensure this hyperplane is between the two classes.

If Σ_1 and Σ_2 are small then $(\Sigma_1 + \Sigma_2)^{-1}$ becomes singular that lead to vanishing the impact of $(\vec{\mu}_2 - \vec{\mu}_1)$, i.e., a solution that leads to singular $(\Sigma_1 + \Sigma_2)^{-1}$ dominates all other solutions, no matter the distance between the centers of the classes [17]. This, however, is not desirable as it is important that the classes centers are apart from one another to be able to distinguish between them. This scenario occurs particularly when the number of instances in a class is smaller than the number of dimensions n. Also, the threshold k is effective only if the distribution of the classes are similar, that might not be the case in many datasets.

One way to extend this algorithm to deal with multiclass classification problems is to use one-vs-all or one-vs-one strategies proposed in [11]. Direct-LDA is another version of LDA that can handle multiple classes.

2.1.6 Direct-LDA

Direct-LDA is a variant of LDA that handles multiple classes [18]. Between-class spread for Direct-LDA is formulated by:

$$S_B = \frac{1}{c}\sum_{i=1}^{c}(\mu_i - \mu)(\mu_i - \mu)^T \quad (1)$$

where μ_i is the average of the instances in the class i, μ is the average of all μ_is. The within-class spread is defined by:

$$S_W = \sum_{i=1}^{c} Cov(X^i) \quad (2)$$

where X^i is a $m_i \times n$ matrix, m_i is the number of instances in the class i, each row is an instance of the class i, and $Cov(.)$ is the covariance operator. In the multiclass case, the optimum value for ω (that is not a vector anymore but a $n \times (c-1)$ matrix) is then the first $c-1$ eigenvectors corresponding to the $c-1$ largest eigenvalues of $S_W^{-1}S_B$. If the number of dimensions is smaller than the number of classes then the algorithm might not find effective ω to distinguish between classes [17]. Another limitation with this formulation is that, if the number of dimensions is smaller than the number of instances in one of the classes,

the covariance matrix for that class becomes non-full rank. In addition, if the distance between two classes is large it may dominate the spread of the classes (S_B) and lead to ineffective transformation [19].

2.1.7 Decision tree (DTR)

Decision tree (DTR) is a tree structure for which each interior node corresponds to one of the input variables and each leaf represents a class label. The outgoing edges from each interior node represent the decision made for variable values in that node. For a given instance, a path from the root of the tree that follows the values of each variable lead to the class label for that instance. The tree is trained (e.g., by the method proposed in [20]) according to the given instances in the training set.

2.2 Derivative-free optimization methods

In this section we provide a brief background information about optimization algorithms we use in this paper.

2.2.1 Quasi-Newton (QN)

The aim of QN is to find a point in a search space that its gradient is zero. The method assumes that the objective function can be estimated by a quadratic function around the local optimum and finds the root of the first derivative of the objective function by generalizing the secant method. We use Broyden-Fletcher-Goldfarb-Shanno (BFGS) [8] in this article to constrain the solutions of the secant equation as this method is frequently used in the literature and provide acceptable practical performance. The finite difference gradient approximation is usually used for objective functions that their gradient cannot be calculated analytically.

2.2.2 Evolutionary strategy (ES)

Evolutionary algorithms work based on a population of candidate solutions that are evolved according to some rules until they converge to an optimum solution. Examples include particle swarm optimization [21] and evolutionary strategy [7], each has its own specific properties that make them advantageous/disadvantageous on various types of problems. The aim of these methods is to use information coded in each individual in the population (with the size λ) and update them to find better solutions. For example, evolutionary strategy (ES) generates new individuals using a normal distribution with the mean of the current location of the individual and an adaptive variance, calculated based on the distribution of "good" solutions. Covariance matrix adaptation evolutionary strategy (CMAES) employs similar idea but updates the covariance matrix of the normal distribution (rather than the variance alone) to generate new instances that accelerates convergence to local optima. This idea takes into account non-separability of dimensions during the optimization process, hence, would be more successful when the variables are interdependent. See [7] for detail of these methods.

3 Proposed algorithm

We define a ψ_κ by a tuple $< F, f, \Omega >$, a surjective transformation F, a discriminator f, and an optimization problem Ω that aims to find the best F such that

$$\forall i \in \{1, ..., c\}, \forall x \in S_i, f(F(\vec{x})) = i$$

where f denotes the class index of the transformed instance $\vec{x}$. Although only F and f are required to classify a set of instances, the optimization problem, Ω, is also very important component to ensure efficiency of the model and the discriminator.

In this paper, we consider that $F : \mathbb{R}^n \rightarrow \mathbb{R}^p$, and $f : \mathbb{R}^p \rightarrow \mathbb{R}$. It is usually assumed that the discriminator function f is constant while the parameters of the transformation F are formulated into Ω and optimized through an optimization procedure. In SVM with linear kernel, for example, the function F is defined by $F(\vec{x}) = \vec{x}\vec{\omega}^T - b$, $\vec{x} \in \mathbb{R}^n$, $f(F(\vec{x})) = sign(F(\vec{x}))$. MLP with no hidden layer assumes that $F(\vec{x}) = t(\vec{x}M_{n\times p} + \vec{b})$ and $f(F(\vec{x})) = F(\vec{x})$, where t is the activation function (usually tan or log sigmoid), and Ω is to minimize the average of $(f(F(\vec{x}^i)) - G(\vec{x}^i))^2$ over all given instances (i) where $G(\vec{x}^i)$ is the class of $\vec{x}^i$.

Let us assume that the instances in each class S_i are random variables that follow a distribution with specific moments. The optimization problem Ω for the optimization of distribution difference (ODD) algorithm is to find a transformation $F : \mathbb{R}^n \rightarrow \mathbb{R}^p$ such that the the gravity center of the transformed instances by F that are in different classes are as far as possible from one another while the distance among transformed instances that are in the same class is minimized. After optimization of F, a discriminator f could be simply defined by the distance between the given instances and the centers of the classes. We formulate F, f, and Ω for ODD in the remaining of this section.

3.1 The optimization problem Ω for ODD

Let $X^{(k)}_{m_k\times n}$ include all given instances of the class k (a subset of S_k), each row corresponds to one instance. We transform each row of this matrix by the function F to form $Y^{(k)}_{m_k\times p}$. We define $\vec{a}^k$, a p dimensional vector, as the center of gravity of all m_k instances in $Y^{(k)}_{m_k\times p}$ as follows:

$$\vec{a}^k = \frac{1}{m_k}\sum_{i=1}^{m_k}\vec{y}_i \tag{3}$$

where $\vec{y}_i$ is the i^{th} row of $Y^{(k)}_{m_k\times p}$. We also define the scalar v^k, as the norm of the eigen values of the covariance matrix of $Y^{(k)}_{m_k\times p}$:

$$v^k = ||Eig(Cov(Y^{(k)}_{m_k\times p}))|| \tag{4}$$

where $Cov(.)$ is the covariance operator and $Eig(.)$ calculates the eigen values of its input matrix. The value of v^k indicates how the instances in the class k have been spread around their center along with their most important directions (eigen vectors). The aim of ODD is to adapt the transformation F such that v^k is minimized for all k while the distances among all possible gravity centers are maximized. This could be formulated by a multiobjective optimization problem:

$$\Omega = \begin{cases} \text{maximize } ||\vec{a}^i - \vec{a}^j|| & \text{for all } j > i \\ \text{minimize } v^i & \text{for all } i \end{cases} \tag{5}$$

where $i, j \in \{1, ..., c\}$. The problem contains c minimization and $\frac{c(c-1)}{2}$ maximization objectives. We use the following remarks to convert this multiobjective problem to a single objective problem:

Remark 1. *Let us assume that $A_i(x)$ is a function and $A_i(x) > 0$ for all i and x. A solution that minimizes $\sum_i A_i(x)$ is on the true Pareto front of the multiobjective optimization problem: "minimize $A_i(x)$ for all i"*

Remark 2. *Let us assume that $A_i(x)$ is a function and $A_i(x) > 0$ for all i and x. A solution that maximizes $\prod_i A_i(x)$ is on the true Pareto front of the optimization problem: "maximize $A_i(x)$ for all i"*

The proofs for both these remarks are elementary and could be done by contradiction.

Using remark 1 and remark 2, the multiobjective optimization problem defined in Eq. 5 can be transformed to a single objective optimization problem defined by:

$$\Omega = \text{minimize } \frac{\gamma + \sum_{k=1}^{c} v^k}{\left(\prod_i^c \prod_{j=i+1}^c ||\vec{a}^i - \vec{a}^j||\right)^{\frac{1}{c(c-1)}}} \quad (6)$$

where γ is a positive constant (set to 1 in our experiments) to ensure that, among all possible solutions for which $\sum_{k=1}^{c} v^k = 0$, the one that maximizes centers distances ($\prod_i^c \prod_{j=i+1}^c ||\vec{a}^i - \vec{a}^j||$) is preferred. Because the growth rate of the denominator is factorial, we have used the regulator $\frac{1}{c(c-1)}$ to balance the growth rate of the nominator and denominator. This ensures balancing the importance of distinct gravity centers while minimizing the spread of instances within each class. The product (rather than a simple weighted summation) enforces the optimizer to find solutions that impose scattered centers for classes. This is an extremely important point as, otherwise, the optimization algorithm may find a solution that maps some of the centers close to one another while move the remaining centers far from the others, that is not desirable.

Note that the transformation of the multiobjective optimization (Eq. 5) to its single objective form (Eq. 6) is effective only if the objectives are assumed to be equally important, that is the case here.

3.2 The discriminator function f for ODD

After solving Ω (Eq. 6), we need to identify to which class a given vector $\vec{y}$ belongs. We define the function $f =< f_1, f_2, ..., f_c >$ as follows:

$$f_k(\vec{y}) = \frac{D_k}{\sum_{j=1}^{c} D_j} \quad (7)$$

where $D_k = ||\vec{y} - \vec{a}^k||$ is the distance between $\vec{y}$ and the center of the class k. The smaller the value of f_k is, the more likely that the vector $\vec{y}$ belongs to the class k. One can calculate $1 - f_k(\vec{y})$ and then normalize the results to get probabilities of $\vec{y} \in S_k$. This leads us to the following formula for f_k:

$$\begin{aligned} f_k(\vec{y}) &= 1 - \frac{1 - \frac{D_k}{\sum_{j=1}^{c} D_j}}{\sum_{i=1}^{c}(1 - \frac{D_k}{\sum_{j=1}^{c} D_j})} \\ &= \frac{1}{c-1} - \frac{D_k}{(c-1)(\sum_{i=1}^{c} D_i)} \end{aligned} \quad (8)$$

Clearly, $f_k(\vec{y}) \in [0, 1]$, that could be interpreted as a measure for the probability of $\vec{y} \in S_k$. In order to convert the results to a categorical value (i.e., converting generative results to the discriminative results), we use the threshold found in the training set that maximizes the area under the curve (AUC) and use that threshold to discriminate the final results in test cases. This thresholding strategy is used for all generative methods in this article (Direct-LDA, ODD, and MLP) unless specified.

3.3 The transformation function F for ODD

We consider the linear case for F in this paper where $F(\vec{x}) = \vec{x} \times M_{n\times p} + \vec{r}$, $\vec{x}$ is a n dimensional vector. Accordingly, Ω in Eq. 6 is dependent only on $M_{n\times p}$ and $\vec{r}$ that introduces $p(n+1)$ variables. In the rest of this paper, we combine $M_{n\times p}$ and $\vec{r}$ and denote $M_{n'\times p}$ where $n' = n + 1$. This, of course, would assume that an instance $\vec{x}^i$ is presented as $< x_1^i, ..., x_n^i, 1 >$. One can see ODD with this setting is very similar to MLP with no hidden layer that uses a different energy function: MLP uses the mean square error of the outputs of instances and expected classes independently while ODD uses the idea of centrality of instances that are in the same class.

Not all problems could be effectively transformed by a linear function. Hence, one may add nonlinear flexibility to F by introducing a function $g : \mathbb{R}^p \to \mathbb{R}^p$ where $F(\vec{x}) = g(\vec{x} \times M_{n'\times p})$. If the function g is nonlinear then the final model could classify instances that are nonlinearly separable. The choice of the function g is problem dependent and could be done through a trial and error procedure. We will test $g(\vec{x}) = tanh(\vec{x})$ in the experiment section as a candidate to introduce non-linearity to our algorithm. This function has been frequently used for this purpose in MLP articles. Note also that, unlike other classification methods that require specific functions for their kernel, the transformation F for ODD could be defined in more generic form as the optimization function Ω for ODD is solved by a derivative-free optimization method.

3.4 Candidate optimization methods for ODD

The optimization problem, Ω, introduced in Eq. 6 is nonlinear. It is also difficult to calculate the gradient and Hessian for this equation, that leaves us with the methods that are either derivative-free or approximate gradient and Hessian. In addition, it is not clear if this optimization problem is unimodal or multimodal, make the solution even more challenging. In this paper, we use three methods to solve this optimization problem: ES, CMAES, and QN. We are not using PSO as it does not take into account the dependency among variables [21] that is required for the purposes of this paper. The first two methods are stochastic and population-based and have a better exploration ability than the last. However, QN could converge to a local optimum faster than other methods [22]. In this section we compare the computational complexity of these methods.

Let $F(\vec{x}) = \vec{x}M_{n\times p}$, $n_v = np$ the number of variables to optimize, $m = kn_v$, where k is a constant. In order to demonstrate the time complexities in practice, we designed a random dataset in which $m = knp$ instances were uniformly randomly sampled in an n dimensional space (n dimensional uniformly sampled instances) and assigned to two classes randomly. We set $p = 2$ in all examples and applied ES, CMAES, and QN to the objective function of

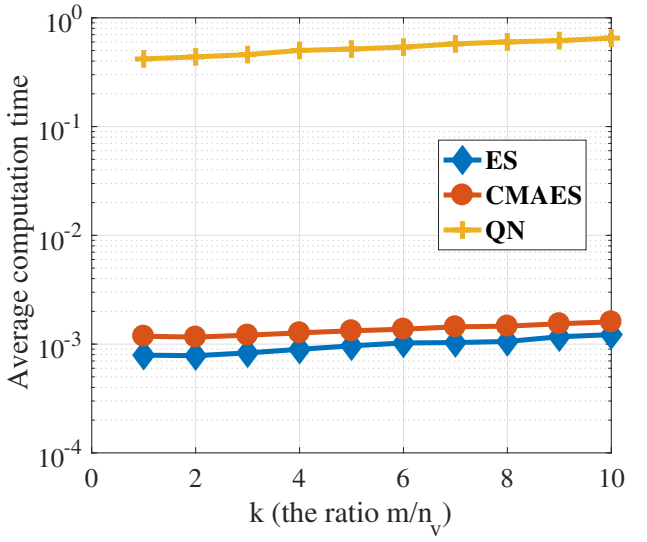

(a)

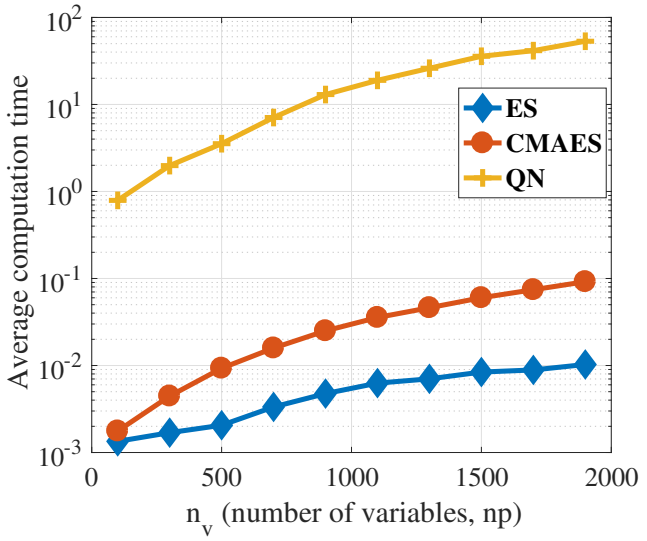

(b)

Fig. 1. The average computation time (in milliseconds) for different methods at each iteration: (a) when k is changed (that changes the number of instances), (b) when n_v is changed.

ODD, each method for 5 iterations, 50 times. The function F was set as specified earlier (linear function), and λ was set to 50 for ES and CMAES. Figure 1(a) shows the results when k was changed from 1 to 10 and $n = 100$. The figure shows that the average computation time for all methods is linear w.r.t k. However, the calculations included much larger constant multipliers for QN that makes the algorithm significantly slower than other methods. Figure 1(b) shows the results when n_v was changed from 100 to 1000 and $m = 2000$. The figure shows that the required time for ES is less than other methods.

Another important factor that contributes to the performance of the optimization method is the convexity of the search space. As there is no reason to assume that the optimization problem defined in Eq. 6 is convex, the exploration ability of the algorithm becomes important. While methods like QN are very efficient in convex spaces, they have difficulty in finding good solutions in non-convex problems. In contrast, ES and CMAES have a better exploration ability that enables them to offer better final solutions in multimodal optimization problems. It is hence beneficial to use a hybrid of ES and CMAES with QN to ensure effective exploration at the beginning of the search while better exploitation at the later stages of the search. Hence, in all of our implementations, we used CMAES in combination with QN for small problems ($n_v \leq 300$) and ES alone for large problems ($n_v > 300$), set experimentally. Although the number of iterations for these methods could be set according to n_v, our experiments showed that 100 for CMAES, 100 for QN, and 500 for ES work efficiently for our test cases. Also, ES and CMAES are terminated if their performance is not improved by at least 0.001 in the last 20 iterations. For QN, the algorithm was terminated if the gradient value was smaller than 1e-8.

3.5 An example

Let us give an example to clarify how ODD works. Assume the dataset presented in Fig. 2[1] is given that contains 140 points in two classes (70 data points in each class) in which each instance $\vec{x}_i$ is 4 dimensional.

We set the transformation F as $F(\vec{x}') = \vec{x}' M_{5\times 2}$, where $\vec{x'}_i =< x_i^1, x_i^2, x_i^3, x_i^4, 1 >$, and $\vec{x}_i^k$ is the k^{th} element of the instance $\vec{x}_i$. We then solve the optimization problem, Ω (Eq. 6) to find the matrix M. After optimization, we found

$$M = \begin{bmatrix} 9.169 & -99.006 & 80.214 & -6.983 & -12.605 \\ 7.111 & -11.865 & -36.473 & 39.129 & -4.548 \end{bmatrix}^T$$

where T is the transpose operator. The transformed dataset before and after optimization of M has been shown in Fig. 3.a and b (filled circles), respectively. The figures also indicate the center of the distributions of transformed instances (crosses) as well as misclassified instances when the value of f was thresholded by an arbitrary value 0.5. With this threshold, the algorithm has classified 133 instances (over all 140) correctly. With an optimized threshold, this could be improved to 137 correctly classified cases.

3.6 Comparison with LDA

In this section we compare ODD with Direct-LDA on some artificially-generated datasets to demonstrate the differences between the methods.

- **Db1**: contains two sets of time series with identical length (300 samples). Both series are in the form of $a\sin(2\pi zt) + 10r\cos(20\pi t) + 10r\cos(400\pi t) + N(0, 30)$, where $N(0, 30)$ generates random values from a normal distribution with mean equal to zero and standard deviation 30. The value of a is picked uniformly randomly from $[15, 25]$ for both series. The frequency z is picked randomly (uniform distribution) from $[65, 75]$ for 50 instances while it is picked randomly (uniform distribution) from $[15, 25]$ for 500 instances. The first 50 instances are labeled as 1 and

1. This dataset is in fact a part of the crab gender dataset [23] where only 4 of the attributes were used for this example. This dataset has been used in its complete form in our experimental results.

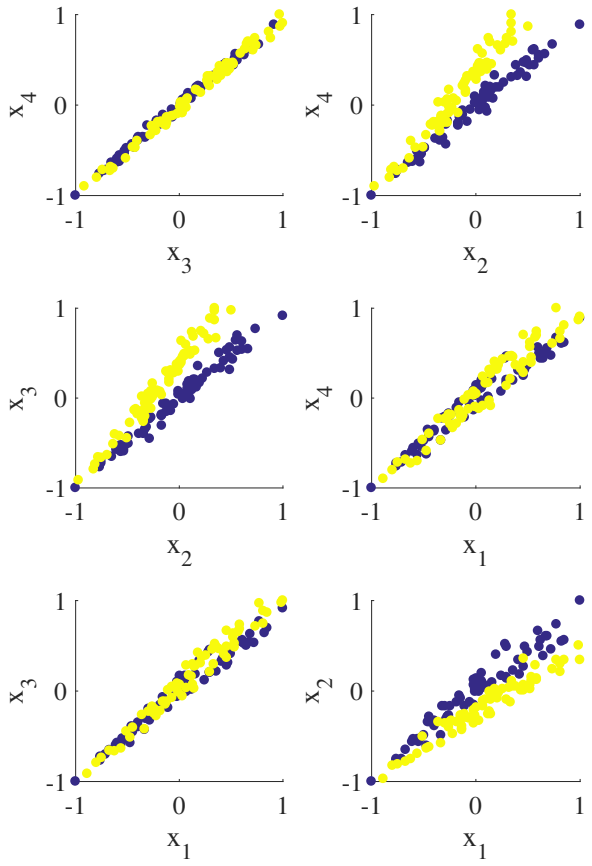

Fig. 2. Illustration of the instances from the Crab gender in four dimensions.

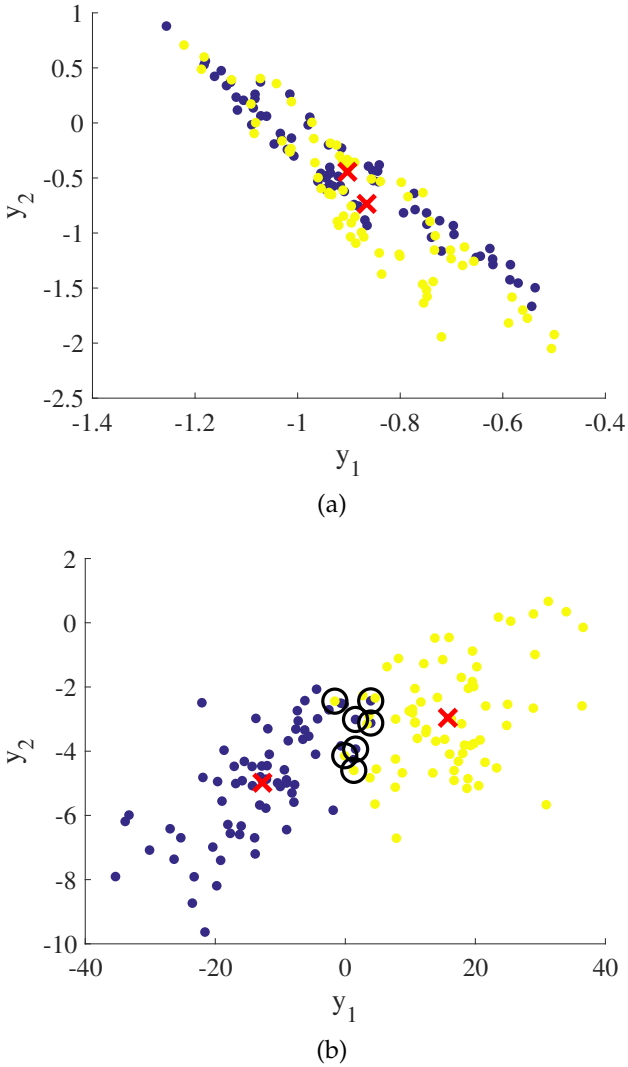

Fig. 3. Application of ODD to the Crab gender dataset (4 dimensions) (a) before, (b) after, optimization of the parameters of F.

the rest are labeled as 2. A similar set is used for testing purposes but with 500 instances from each class.

- **Db2**: similar to the first dataset, but this time we have 4 different frequency ranges: $[10, 20], [30, 40], [50, 60], [70, 80]$ in 4 classes. We place 50, 250, 500, and 10 instances from each time series for training and the same number of instances for testing.
- **Db3**: two sets of time series, 100 instances in the first and 1000 in the second. Each sample of each series is generated by $N(a, b)$. For the time series in the first group it is either $a = 40$ and $b \in [95, 105]$ (uniform distribution) OR $b = 80$ and $a \in [49, 51]$ (uniform distribution). For the time series in the second group it is either $a = 40$ and $b \in [55, 65]$ (uniform distribution) OR b=80 and $a \in [29, 31]$ (uniform distribution). It is clear that these two sets are not distinguishable by their variance or mean, but both at the same time.
- **Db4**: was taken from "Figure 3" of [24]. It includes 788 instances in 7 classes in 2 dimensions[2].

2. The data is available online at https://cs.joensuu.fi/sipu/datasets/.

We applied ODD and Direct-LDA to these four datasets, results have been reported in Fig. 4. For both methods we used the distance between the instances and the closest distribution center without any thresholding (function f, see section 3.2). Clearly, ODD outperforms Direct-LDA in all training sets. The most important denominator in all of these sets was that either the number of dimensions was smaller than the number of classes (e.g., Db4) or the number of instances in at least one of the classes was smaller than the number of dimensions (e.g., Db3). Both of these scenarios cause Direct-LDA to fail (the method could not find any solution for Db2 and Db4) while ODD performs fine in these scenarios. Note that the second scenario is very common in time series, i.e., large number of samples, each represent one dimension, and small number of instances for each class.

4 Experimental Results

In this section we provide comparative results of the algorithm with other state-of-the-art classification methods. We apply ODD together with other methods to 12 standard classification problems.

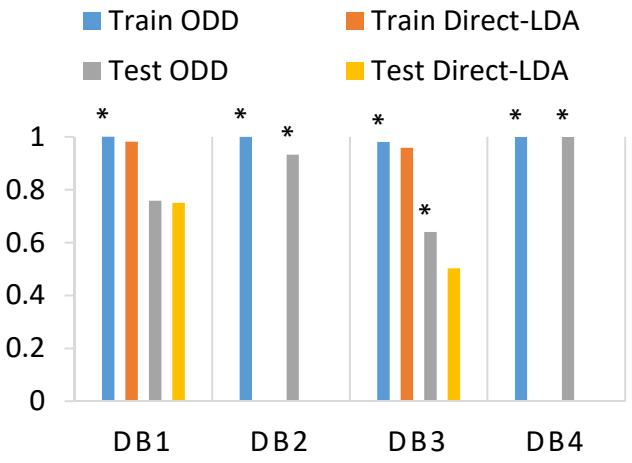

Fig. 4. The results of the accuracy (area under the curve) of Direct-LDA in comparison to ODD. Star indicates the significance of the comparison (0.05 confidence, t-test). Direct-LDA could not find any solution for Db2 and Db4.

4.1 Comparisons procedure

In this subsection we introduce the datasets, pre-processes, and algorithms settings used for comparisons.

4.1.1 Datasets

We use 12 datasets for comparison among different classifiers, namely, Breast cancer (BC), Crab gender (GD), Glass chemical (GC), Parkinson (PR) [25], Seizure detection (SD) [26], Iris (IR), Italian wines (IW), Thyroid function (TF), Yeast dataset (YD) [27], Red wine quality (RQ) [28], White wine quality (WQ) [28], and Handwritten digits (HD). Main characteristics of these datasets have been reported in Table 1. These datasets have been used frequently in previous classification studies as standard benchmarks. We used the one-vs-rest presentation for the classes, hence, the class of each instance was represented by a binary vector with the length c.

The SD dataset includes interacornial electroencephalogram (iEEG) of 12 patients (4 dogs and 8 Human) with variable number of channels (Table 2 shows the details of this dataset) [26]. There are 2 classes, namely seizure (ictal) and no-seizure (interictal), in the dataset with various number of instances and iEEG channels for each patient. While each seizure event might take up to 60 seconds, each instance of ictal or interictal in the dataset indicates only 1 second of an event from all iEEG channels. As the properties of the signals that belong to the same ictal event could be similar, including different segments of a single event in both training and testing sets may simplify the problem. Hence, we used all ictal segments that belonged to the same seizure event in either testing or training set, but not both. The segment index is available with the dataset that could be used to reconstruct signals from the same event. This procedure is usually used for cross-validation in seizure detection and prediction literature [26].

4.1.2 Pre-processes and performance measure

We preprocessed the instances in the SD dataset by calculating the fast Fourier transform of each channel and concatenating these transformed signals to generate one large signal (FFT of the channels one after another). The length of this signal is a function of the number of channels of the iEEG device. We used frequencies from 1 to 50 Hz only, as it showed sufficiently accurate presentation of seizure [26]. For subject 1, for example, the preprocessed signal was 16×49 samples long (16 channels, 1 Hz to 50 Hz FFT). The classifiers were trained for each patient independently. The training set was generated by selecting 70% of the seizure events randomly and the rest of the seizure events were left for test. We also selected 70% of non-seizure signals randomly for training and left the rest for test. For each selected set of training signals we trained KNN, MLP, SVM, DTR, LDA, and ODD models and tested their performances on the remaining instances. This was done for 50 independent runs to decrease the impact of possible biased selection in train and test sets. As the number of instances in each class

TABLE 1
The datasets used for comparison purposes in this paper. n is the number of variables and c is the number of classes in each dataset. The number of instances in each class has been reported in the last column.

Dataset name	n	c	# instance in each class
BC	9	2	$< 458, 241 >$
CG	6	2	$< 100, 100 >$
GC	9	2	$< 51, 163 >$
PR	22	2	$< 48, 147 >$
SD*	?	2	?
IR	4	3	$< 50, 50, 50 >$
IW	13	3	$< 59, 71, 48 >$
TF	21	3	$< 166, 368, 6666 >$
YD	8	10	$< 463, 5, 35, 44, 51, 163, 244, 429, 20, 30 >$
RQ	11	6	$< 10, 53, 681, 638, 199, 18 >$
WQ	11	7	$< 20, 163, 1457, 2198, 880, 175, 5 >$
HD	784	10	$< 6903, 7877, 6990, 7141, 6824, 6313, 6876, 7293, 6825, 6958 >$

*The seizure detection dataset includes 12 patients, each of them has their own number of variables and instances in different classes. See Table 2.

TABLE 2
Details of the seizure detection (SD) dataset. The first value in the last column is the number of instances of seizure (one second each) and the second number is the number of instances of non-seizure (one second each). The seizure instances are one-second segments from different seizure events.

Patient	# of Channels	Sampling rate (Hz)	# instance in each class
Subject 1	16	400	$< 178, 418 >$
Subject 2	16	400	$< 172, 1148 >$
Subject 3	16	400	$< 480, 4760 >$
Subject 4	16	400	$< 257, 2790 >$
Subject 5	68	500	$< 70, 104 >$
Subject 6	16	5000	$< 151, 2990 >$
Subject 7	55	5000	$< 327, 714 >$
Subject 8	72	5000	$< 20, 190 >$
Subject 9	64	5000	$< 135, 2610 >$
Subject 10	30	5000	$< 225, 2772 >$
Subject 11	36	5000	$< 282, 3239 >$
Subject 12	16	5000	$< 180, 1710 >$

for SD dataset is imbalance (the ratio ictal to interictal is almost 2:19 in average), this dataset forms a good test for the models.

For other datasets, the variables were normalized to $[-1, 1]$ for the instances in the train set. The same mapping that normalized the training set was then applied to the test set to ensure both train and test sets are in the same domain. The variables with the variance equal to zero across the whole dataset were completely removed from further calculations. For HD dataset, we also used principle component analysis [29] to reduce the number of dimensions to 50, prior to normalization.

For each dataset (except for HD), we randomly selected 70% of instances in each class for training purposes and left the rest for tests. After this selection, all methods were applied to the training set and then their performances on both training set and test set were measured. The selection and modeling were performed for 50 independent runs to ensure the results are not biased towards a specific combination of the instances in the train set. For each tune, the training set was remained unchanged for all methods. For HD dataset (70,000 instances), the number of selected instances for training purposes at each run was set to 5% (rather than 70%) from each class and the remaining instances were left for testing purposes.

We used area under the curve (AUC) of the recipient operation curve (ROC) [30] as a performance measure. For multiclass problems, the overall performance was calculated by averaging AUC for each class separately. As MLP and ODD are generative classifiers, we used the best threshold found on the ROC curve of the training set to threshold the results of the test set and convert the results to discriminative. The same threshold was used for testing purposes.

We used the t-test (confidence 0.05) for statistical comparisons between the performance of each method and ODD.

4.1.3 Algorithms settings

We compare the results of ODD with NBY, SVM (we used one-vs-all strategy to enable SVM to deal with multiple classes), MLP (with the number of neurons in the hidden layer equal to p (the dimension of the transformation F in ODD) with Levenberg-Marquardt for backpropagation of weights errors), LDA (we used one-vs-all strategy to enable LDA to deal with multiple classes), DTR, and KNN (5 neighbors model). We used Matlab 2016b for implementations and tests. For SVM and LDA we used linear kernels. We tested three settings for ODD that are ODD_1 for which $p = 1$ and $F(\vec{x}) = \vec{x}M_{n' \times p}$, ODD_l for which $p = c$ and $F(\vec{x}) = \vec{x}M_{n' \times p}$, and ODD_n for which $p = c$ and $F(\vec{x}) = tanh(\vec{x}M_{n' \times p})$. Note that all derivatives of ODD could handle multiclass classification with no need for external strategies.

Stopping criteria for MLP was set to gradient $<$ 1e-8 or 1000 iterations. For ODD_1 and ODD_l, the stopping criteria was 100 iterations (constant) of CMAES and then 100 iterations of QN to ensure efficient exploration and exploitation. For CMAES, if the performance was not improving then the algorithm was terminated (improvement for the last 20 iterations was smaller than 0.0001). For QN, if gradient was smaller than 1e-8 then the algorithm was terminated. For large problems in these tests (HD and SD), we only used ES for 500 iterations.

4.2 Comparison with existing classifiers

Table 3 shows comparative results[3] of tested methods for all datasets except for SD. The value in the row i column j shows $P_{i,j} - G_{i,j}$ where $P_{i,j}$ is the number of datasets for which the ODD type indicated in column j performs significantly (based on t-test, confidence 0.05) better than the method indicated in the row i and $G_{i,j}$ is the number of datasets for which the method indicated in the row i performs significantly better than the ODD type indicated in column j, for the performance measure indicated in the "Measure" column. For example, the value 3 in the row 3 (LDA), column 2 (ODD_n) for the measure "Test" indicates that the number of datasets for which ODD_n performs significantly better than LDA is 3 datasets (over 11) more than number of datasets for which LDA performs significantly better than ODD_n.

Clearly the running time of all derivatives of ODD is significantly longer than the running time of NBY, SVM, LDA, DTR, and KNN. In comparison to MLP, however, ODD_l requires significantly less time in majority of datasets. This was not as good for ODD_n mainly because of the overhead of the non-linear function. For ODD_1, this running time was significantly less than MLP's in majority of datasets.

The performance of ODD_l was significantly better than all other methods in majority of datasets in the test sets (all values in that column are positive). The algorithm also performed significantly better than other methods in majority of datasets in the training set except in comparison to MLP and DTR.

The performance of ODD_1 was significantly better than all other methods in majority of datasets for test sets except in comparison to MLP and LDA. For training sets, ODD_1 performs significantly better than KNN, LDA, SVM, and NBY in majority of datasets while DTR and MLP perform better than ODD_1 in majority of datasets for training set.

The performance of ODD_n was significantly better than all other methods in majority of datasets in the test sets except in comparison to MLP, where there is a draw (for 5 dataset MLP performs better while for 5 ODD_n performs better). ODD_n also performed significantly better than other methods in the training set in majority of datasets except in comparison to MLP and DTR.

4.3 SD dataset results

Table 4[4] shows comparative results of tested methods for SD dataset. ODD_l and ODD_1 performed significantly better than all other methods in majority of subjects for the test set. For the training set, however, ODD_l and ODD_1 performed better than all methods except LDA and SVM for majority of subjects. ODD_n performed better than all other methods for training set for majority of subjects except in comparison to SVM, LDA, and MLP. For test sets, the algorithm performs significantly better than all other methods in majority of subjects except in comparison to LDA.

3. Details of this experiment is available in Appendix A.
4. Details of this experiments are available in Appendix B

TABLE 3
Comparison results among different classification algorithms. Each row indicates the comparative results with one type of ODD in terms of time, training set performance, and test set performance. Positive values indicate better performance (in terms of different measures) of the ODD types.

	ODD_l			ODD_n			ODD_1		
Measure	Time	Train	Test	Time	Train	Test	Time	Train	Test
NBY	-11	11	7	-11	10	4	-11	7	5
SVM	-11	11	8	-11	11	7	-11	7	5
LDA	-11	9	2	-11	9	3	-11	5	-1
MLP	7	-11	1	1	-11	0	9	-11	-2
DTR	-11	-5	7	-11	-4	7	-11	-7	5
KNN	-11	6	3	-11	7	3	-11	3	1

TABLE 4
Comparative results among different classification algorithms when they were applied to SD dataset. The values in the table are similar to what was discussed for Table 3.

	ODD_l			ODD_n			ODD_1		
Measure	Time	Train	Test	Time	Train	Test	Time	Train	Test
NBY	-12	12	11	-12	12	10	-12	12	11
SVM	-12	-6	12	-12	-9	12	-12	-6	12
LDA	-12	-1	6	-12	-5	-1	-12	0	5
MLP	-12	3	12	-12	-1	11	-12	3	12
DTR	-12	10	10	-12	6	9	-12	10	10
KNN	-12	12	12	-12	12	10	-12	12	12

In terms of running time, all types of ODD performed significantly slower than other method.

4.4 Sensitivity to between-class imbalance datasets

ODD is not sensitive to the number of instances in each class. The reason is that ODD optimizes characteristics (mean and spread) of a class that is independent of the number of instances in that class. It is easy to see that, assuming instances in each class represent the exact mean and spread of that class, the performance of ODD is independent of the number of instances in the classes. Although in reality the instances in each class might not reflect the exact mean and spread of the class, it is expected that these parameters are preserved in the given data to some extent to enable the algorithm to perform well. Similar assumption is required by any other classification algorithm (i.e., training set must reflect the characteristics of the data in each class), otherwise, it is impossible to design any effective classifier.

We test the between-class imbalance sensitivity of ODD by using the CG dataset. We divide this dataset to two subsets, one for training and the other for test, with 70% of instances from the first class were always present in the train set. The ratio of the instances from the second class was $100r\%$, $r \in \{0.1, 0.2, ..., 0.7\}$. As the number of instances in each class is equal to 100 in this dataset, the changes in r causes changes in the balance of the number of instances in the training set. We run NBY, SVM, LDA, DCT, KNN, and ODD for training and predict the class of the remaining instances. The average of the performance (area under the curve, AUC, was used for performance measure) of the methods over 50 runs has been reported in Fig. 5. For ODD, p was set to 1 to ensure a fair comparison with SVM, hence,

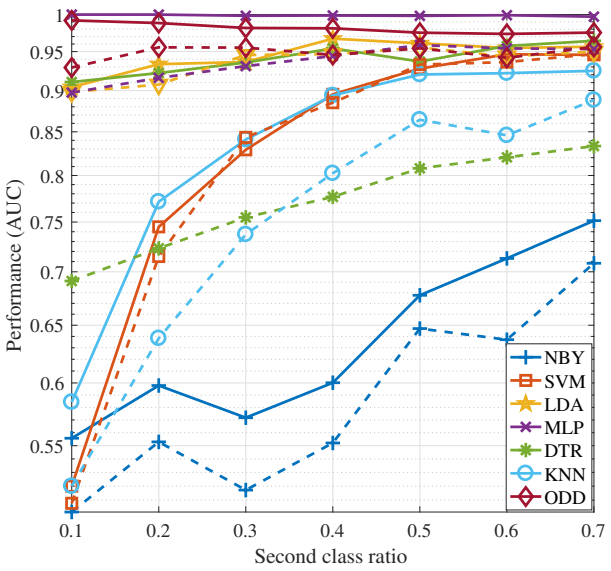

Fig. 5. Comparative results in terms of the sensitivity to the number of instances in each class. The dashed lines are the test results while the solid lines are the training results. ODD and MLP are affected only slightly by imbalanced between-class number of instances. The CG dataset was used for this test.

$F(\vec{x}) = \vec{x}M_{n' \times 1}$ where $n' = 7$ as the number of variables in the dataset is 6.

Figure 5 that the performance of ODD and MLP changed slightly when the class ratio was changed from 0.1 to 0.7. This, however, is not correct for SVM, NBY, LDA, DCT, and KNN and these methods have been affected significantly by changing the imbalance ratio.

4.5 Overfitting

Over-fitting is a very common issue among classification methods. It is usually measured by the extent the performance of a classification method could be generalized to unseen instances. To formulate overfitting for a classification method, we calculate the performance extension index defined by:

$$G_{indx} = \frac{\text{Test performance}}{\text{Train performance}} \times \text{Test performance} \quad (9)$$

This index indicates to what extent the performance of the classifier on the train data could be generalized to the unseen data (the fraction) while also takes into account how well the algorithm performs in the testing set (the second term). The second term has been used in this formula to penalize algorithms that have a very low performance (e.g., 50 percent accuracy in both test and train) on both training set while their bad performance is extendable.

The larger the value of G_{indx} is, the better the algorithm can generalize its performance to the training set, taking into account how good is the algorithm performance in the test set. We calculated the G_{indx} for all methods based on their results in section 4.2. Figure 6 shows the average ranking of the algorithms based on their G_{indx} value over all 11 datasets (the smaller the ranking the better). The figure indicates that ODD_l has the best ranking among these

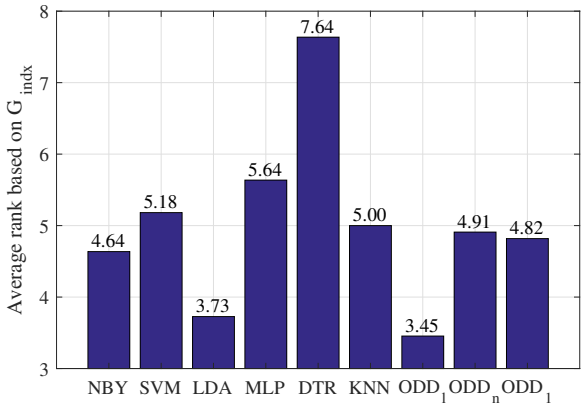

Fig. 6. Comparison results among different types of ODD and other classifiers in terms of G_{indx}. The vertical axis indicates the average rank of the methods in terms of their G_{indx}. The smaller the rank the more successful the algorithm was in generalization of its performance according to the G_{indx} measure.

methods in terms of its G_{indx}. LDA is in the second place with a small margin and NBY is in the third place. Also, ODD_n has better generalization ability comparing to all non-linear classification methods tested in this article (i.e., MLP, DTR, and KNN).

5 Conclusion and Future Works

This paper introduced a new method based on optimization of distribution differences (ODD) for classification. The aim of the algorithm was to find a transformation to minimize the distance between instances in the same class while maximize the distance between gravity centers of the classes. The definitions of ODD allows the use of any transformation, however, this paper only tested linear transformation and a particular non-linear transformation. The algorithm was applied to 12 benchmark classification problems and results were compared with state-of-the-art classification algorithms. Comparisons showed that the proposed algorithm outperforms previous methods in most datasets. Results also showed that the method is less sensitive to the between class imbalance and has a better ability to maintain its performance for the instances that have not been used in the training process.

This paper only experimented with reducing the number of dimensions to a value equal to the number of classes. However, ODD design allows flexible dimensionality reduction/increase that could be beneficial in some datasets. Hence, as a future work, we are going to design an adaptive method to increase the number of dimensions in the linear transformation (p) gradually to find the best size for the transformation matrix. A small transformation matrix could be beneficial especially for hardware implementations. In addition, the method could be extended to make use of a multiple layers in which each layer reduces the number of dimensions that could improve the performance of the algorithm for some classification tasks. Another interesting direction would be to change the stopping criteria and test its impact on overfitting, i.e., stop the algorithm as soon as the instances were separable in different classes.

Appendix-A

Table 5 shows the average results of 50 independent runs for all algorithms. The values in the table have been prefixed by three characters. The character in position one, two, and three after each value at a specific row and column indicate the results of the statistical test (t-test with confidence 0.05) between the method indicated in that column and ODD_l, ODD_n, and ODD_1, respectively, for the dataset at that row. "*", "-", and "+" at a position indicate that that result is statistically worst, the same, or better than ODD_l, ODD_n, and ODD_1, depending on the position of the character. For example, the value "96.37^{*-*}" in the row "BC", measure "Test", column NBY indicates that the average performance (AUC) of the method NBY was 96.37 on the training set, that was significantly worse than ODD_l and ODD_1 while statistically the same comparing to ODD_n

Appendix-B

Table 6 shows the average results of 50 independent runs for all algorithms when they were applied to SD. The values in the table have been postfixed by three characters that have the same definitions as in Table 5.

References

[1] A. Y. Ng and M. I. Jordan, "On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes," *Advances in neural information processing systems*, vol. 2, pp. 841–848, 2002.

[2] S. Haykin and N. Network, "A comprehensive foundation," *Neural Networks*, vol. 2, no. 2004, p. 41, 2004.

[3] J. A. Suykens and J. Vandewalle, "Least squares support vector machine classifiers," *Neural processing letters*, vol. 9, no. 3, pp. 293–300, 1999.

[4] G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew, "Extreme learning machine: theory and applications," *Neurocomputing*, vol. 70, no. 1, pp. 489–501, 2006.

[5] P. Branco, L. Torgo, and R. P. Ribeiro, "A survey of predictive modeling on imbalanced domains," *ACM Computing Surveys (CSUR)*, vol. 49, no. 2, p. 31, 2016.

[6] D. M. Hawkins, "The problem of overfitting," *Journal of Chemical Information and Computer Sciences*, vol. 44, no. 1, pp. 1–12, 2004, pMID: 14741005.

[7] H.-G. Beyer and H.-P. Schwefel, "Evolution strategies–a comprehensive introduction," *Natural computing*, vol. 1, no. 1, pp. 3–52, 2002.

[8] R. Fletcher, *Practical methods of optimization*. John Wiley & Sons, 2013.

[9] N. S. Altman, "An introduction to kernel and nearest-neighbor nonparametric regression," *The American Statistician*, vol. 46, no. 3, pp. 175–185, 1992.

[10] D. J. Hand and K. Yu, "Idiot's bayesnot so stupid after all?" *International statistical review*, vol. 69, no. 3, pp. 385–398, 2001.

[11] C.-W. Hsu and C.-J. Lin, "A comparison of methods for multiclass support vector machines," *IEEE Transactions on Neural Networks*, vol. 13, no. 2, pp. 415–425, Mar 2002.

[12] D. W. Ruck, S. K. Rogers, M. Kabrisky, M. E. Oxley, and B. W. Suter, "The multilayer perceptron as an approximation to a bayes optimal discriminant function," *IEEE Transactions on Neural Networks*, vol. 1, no. 4, pp. 296–298, 1990.

[13] K. Levenberg, "A method for the solution of certain non-linear problems in least squares," *Quarterly of applied mathematics*, vol. 2, no. 2, pp. 164–168, 1944.

[14] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters," *Journal of the society for Industrial and Applied Mathematics*, vol. 11, no. 2, pp. 431–441, 1963.

[15] S. K. Pal and S. Mitra, "Multilayer perceptron, fuzzy sets, and classification," *IEEE Transactions on neural networks*, vol. 3, no. 5, pp. 683–697, 1992.

TABLE 5
Comparison results among 6 well-known classification methods with ODD. ODD_l refers to a setting where $p = c$, c is the number of classes, and a linear kernel is used for F. ODD_n uses the same value for p with $F(\vec{x}) = tanh(\vec{x}M_{n' \times p})$. ODD_1 uses a linear F while sets $p = 1$. The values in each row is the average (over 50 runs) of results. The row "Time" reports the average of time in milliseconds over 50 runs.

Dataset	Measure	**NBY**	**SVM**	**LDA**	**MLP**	**DTR**	**KNN**	ODD_l	ODD_n	ODD_1
BC	Time	20.4	12.6	10.5	991.9	8.8	8.4	99	423.7	63.2
	Train	96.37***	96.9***	95.26***	**98.63**+++	97.96+-+	97.55***	97.77	97.91	97.77
	Test	96.37*-*	96.49*-*	95.13***	95.42***	93.98***	96.71—	**97.06**	96.6	97.06
CG	Time	6.7	5.4	5.3	154.3	5	4.6	148.2	343.6	98.2
	Train	76.39***	94.86***	95.79***	**99.66**+++	96.16***	93.03***	96.7	97.46	96.93
	Test	69.93***	94.13—	**94.73**+++	94.33—	82.07***	86.93***	94.08	94.07	94.25
GC	Time	7.9	5.2	5.2	527.8	4.8	4.5	124.1	366.1	84.2
	Train	87.61***	89.86***	89.58***	**98.85**+++	96.93+++	93.08***	95.35	96.07	95.36
	Test	86.58***	87.71***	85.7***	86.74***	88.73***	87.04***	90.44	**90.63**	90.21
PR	Time	13.4	5.3	5.5	720.5	5.1	4.6	325.3	843	165.6
	Train	78.32***	79.42***	84.99***	**97.52**+++	96+++	93.05+-+	88.03	93.86	87.6
	Test	76.64—	74.83-*-	77.64–+	75.83-*-	78.24–+	**85.56**+++	76.29	78.35	75.12
IR	Time	6.2	9.6	6	255.2	4.7	4.6	138.9	421.2	86.1
	Train	97.06***	80.7***	98.53***	**99.68**+++	98.47***	97.64***	98.69	99.32	98.69
	Test	96.43—	78.73***	**98.2**+++	95.7***	96.2–*	97.03—	96.88	96.59	96.86
IW	Time	10.4	9.9	5.9	83.7	4.9	4.1	316.2	719.8	136.4
	Train	98.94**+	99.31**+	99.99+++	**100**+++	98.85**+	98.26**+	99.78	99.76	97.33
	Test	97.94-++	98.04-++	**98.88**+++	95.58**+	91.88***	97.48-++	97.56	96.83	93.83
TF	Time	0	229.2	10.8	8431.5	8.4	5.2	685.1	4261.6	275.6
	Train	N/A***	64.35***	62.57***	98.3+++	**99.55**+++	70.03***	78.4	83.49	76.92
	Test	N/A***	63.8***	61.99***	96.5+++	**98.52**+++	64.83***	76.71	81.72	75.5
YD	Time	0	100.7	18.3	15303.7	11.1	5.6	2601.5	6299.7	454.5
	Train	N/A***	64.07***	76.74**+	**85.26**+++	82.52+++	80.27–+	80.19	80.43	73.87
	Test	N/A***	63.61***	76.23+-+	**77.69**+++	67.1***	75.34-*+	75.26	75.92	69.71
RW	Time	13.1	83.4	10.2	6335.4	13.5	4.3	1386.2	3871.9	395.9
	Train	64.82***	57.58***	62.13***	79.36+++	**82.41**+++	66.3***	77.75	78.06	75.62
	Test	60.57***	57.37***	60.42***	70.67-+-	60.16***	58.92***	**70.71**	68.68	69.92
WW	Time	19	673.6	13.5	21432.9	35.9	4.8	1743.5	7361.5	547.3
	Train	61.74***	52.08***	58.53***	75.6+++	**80.14**+++	67.58***	71.8	70.67	68.85
	Test	59.33***	52.01***	57.66***	**71.05**+++	60.78***	59.18***	65.8	61.99	65.26
HD	Time	92	2099.2	33.1	3222.5	127.9	7.7	20570.5	28141.2	941.2
	Train	93.1*-+	90.63**+	92.73**+	96.39+++	95.95+++	**96.89**+++	93.26	93.13	78.08
	Test	92.34-++	89.23**+	92.06**+	92.46-++	85.11**+	**95.1**+++	92.33	92.19	77.04

[16] R. A. Fisher, "The use of multiple measurements in taxonomic problems," *Annals of eugenics*, vol. 7, no. 2, pp. 179–188, 1936.

[17] H. Gao and J. W. Davis, "Why direct lda is not equivalent to lda," *Pattern Recognition*, vol. 39, no. 5, pp. 1002–1006, 2006.

[18] H. Yu and J. Yang, "A direct lda algorithm for high-dimensional datawith application to face recognition," *Pattern recognition*, vol. 34, no. 10, pp. 2067–2070, 2001.

[19] T. Li, S. Zhu, and M. Ogihara, "Using discriminant analysis for multi-class classification: an experimental investigation," *Knowledge and information systems*, vol. 10, no. 4, pp. 453–472, 2006.

[20] J. R. Quinlan, "Induction of decision trees," *Machine learning*, vol. 1, no. 1, pp. 81–106, 1986.

[21] M. R. Bonyadi and Z. Michalewicz, "Particle swarm optimization for single objective continuous space problems: a review," *Evolutionary computation*, 2016.

[22] I. Loshchilov, "Lm-cma: An alternative to l-bfgs for large-scale black box optimization," *Evolutionary computation*, 2015.

[23] A. Asuncion and D. Newman, "Uci machine learning repository," 2007.

[24] A. Gionis, H. Mannila, and P. Tsaparas, "Clustering aggregation," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 1, no. 1, p. 4, 2007.

[25] M. A. Little, P. E. McSharry, E. J. Hunter, J. Spielman, L. O. Ramig *et al.*, "Suitability of dysphonia measurements for telemonitoring of parkinson's disease," *IEEE transactions on biomedical engineering*, vol. 56, no. 4, pp. 1015–1022, 2009.

[26] A. Temko, A. Sarkar, and G. Lightbody, "Detection of seizures in intracranial eeg: Upenn and mayo clinic's seizure detection challenge," in *Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE*. IEEE, 2015, pp. 6582–6585.

[27] P. Horton and K. Nakai, "A probabilistic classification system for predicting the cellular localization sites of proteins." in *Ismb*, vol. 4, 1996, pp. 109–115.

[28] P. Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis, "Modeling wine preferences by data mining from physicochemical properties," *Decision Support Systems*, vol. 47, no. 4, pp. 547–553, 2009.

[29] H. Hotelling, "Analysis of a complex of statistical variables into principal components." *Journal of educational psychology*, vol. 24, no. 6, p. 417, 1933.

[30] J. A. Hanley and B. J. McNeil, "The meaning and use of the area under a receiver operating characteristic (roc) curve." *Radiology*, vol. 143, no. 1, pp. 29–36, 1982.

TABLE 6
Results for seizure detection dataset.

Dataset	Measure	**NBY**	**SVM**	**MLP**	**LDA**	**DTR**	**KNN**	ODD_l	ODD_n	ODD_1
Subject 1	Time	371.2	24.1	1924.4	94.8	68.1	7.6	6730	7652.3	3354.1
	Train	83.73***	**100**+++	100+++	100+++	99.72***	97.35***	100	99.83	100
	Test	87.11***	95.15***	96.72*-*	82.81***	89.77***	90.72***	97.2	96.76	**97.23**
Subject 2	Time	410.5	48.2	3087.3	170.1	302.2	5	10848.9	13009.4	6720.9
	Train	90.58***	**100**—	100—	100—	99***	93.6***	100	100	100
	Test	69.65***	79.44***	**94.32**+++	77.54***	72.36***	87.98***	92.1	91.78	91.93
Subject 3	Time	827.9	791.2	8542.1	452.5	2258.6	7.5	47136.3	52129.2	33319.5
	Train	91.53***	**100**+++	99.88++-	96.47***	98.78***	96.54***	99.81	99.38	99.83
	Test	91.07***	89.44***	95.3*-*	91***	89.06***	89.43***	**96.77**	95.54	96.76
Subject 4	Time	545.4	689.2	5606.9	292.3	1198.8	5.9	24416.3	28656.7	18166.7
	Train	67.41***	**100**+++	100+++	95.35***	98.53***	96.09***	99.79	99.65	99.81
	Test	83.5***	77.13***	73.07***	80.52***	76.39***	82.61***	**95.33**	91.32	94.47
Subject 5	Time	1532.9	19.6	2163.2	102.5	53.4	4.7	9025.9	9108.6	5912.7
	Train	93.56***	**100**—	99.01—	100—	100—	98.32***	100	100	100
	Test	86.85***	84.09***	**96.52**—	79.54***	79.91***	87.01***	94.74	96.43	94.66
Subject 6	Time	636.6	192.5	5614.7	294.8	680.4	6.1	25360.2	29460.1	18817.2
	Train	92.65***	**100**+++	99.81*+*	97.9***	97.65***	94.34***	99.97	99.6	99.97
	Test	94.9***	95.03***	98.7***	96.7***	90.53***	91.02***	99.34	99.07	**99.35**
Subject 7	Time	1413	388.7	5461.3	600.4	367.2	5.9	31242.9	32758.9	25726.8
	Train	82.76***	**100**-+-	98.16—	100-+-	99.71*-*	90.92***	100	99.67	100
	Test	60.84***	61.44***	**73.55**+-+	60.87***	71.98+-+	55.91***	69.13	72.69	67.12
Subject 8	Time	1621.4	15.1	1619.4	127.2	42.9	4.7	11072.7	11317.2	7080.7
	Train	99.47***	**100**-+-	99.04—	100-+-	100-+-	97.5***	100	99.86	100
	Test	52***	53.75***	64.31*+*	56.21*-*	46.44***	57.78*+*	**71.49**	55.47	71.36
Subject 9	Time	2234.3	1060	12547.6	7992.8	1991.4	8.9	97547	102059.8	76298.9
	Train	80.29***	**100**-+-	99.98-+-	100-+-	99.96*+*	94.44***	100	99.65	100
	Test	91.81-+-	85***	89.66*-*	68.11***	77.67***	86.11***	**91.84**	89.34	91.79
Subject 10	Time	1038.9	335	7782.7	1103.1	573.7	7	51889.3	55736.9	37255.1
	Train	96.21***	**100**+-+	100+-+	100+-+	99.76***	90.89***	100	99.99	100
	Test	85.87***	85.84***	93.31*+*	86.96***	89.29*+*	75***	**95.55**	88.52	95.4
Subject 11	Time	1592.4	808.9	9661.8	1988.7	2235.8	8.1	70730.1	74210.7	52923.8
	Train	91.04***	**100**-+-	84.12***	99.98*+*	99.42***	96.41***	100	99.96	100
	Test	98.6***	99.03***	83.43***	98.38***	96.14***	99.21***	**99.92**	99.85	99.91
Subject 12	Time	447.5	111	3865.3	205	256.1	5.2	16811.3	19884.6	9818.7
	Train	93.85***	**100**+++	99.74*+*	98.89***	98.33***	97.22***	99.98	99.05	99.98
	Test	90.04***	91.77***	**97.95**-++	88.63***	88.09***	93.06***	96.85	95.38	96.82

Because of the size of these problems, we used ES with 500 iterations only.